



# DESIGNING DECISION RIGHTS FOR AI

Ina M. Sebastian, *Research Scientist*

Peter Weill, *Chairman and Senior Research Scientist*

MIT Center for Information Systems Research (CISR)

Thomas Haskamp, *Assistant Professor, Department of Information Systems, University of Münster and Academic Research Fellow, MIT CISR*

Jan vom Brocke, *Professor and Chair of Information Systems and Business Process Management, University of Münster and Academic Research Fellow, MIT CISR*

One of the biggest challenges companies face today is how to design decision rights between humans and artificial intelligence (AI). Should AI be used in hiring, and if so, for which tasks? Should it run marketing campaigns end to end, or augment human judgment? How can companies in regulated industries deploy AI agents<sup>1</sup> without compromising trust?

Drawing on thirty executive interviews conducted in 2025 and 2026,<sup>2</sup> we developed the AI Decision Matrix, a framework that helps leaders determine how humans and AI that can complete tasks autonomously should share decisions based on ambiguity and risk. We use the example of One New Zealand Group Ltd (One NZ), a telecommunications provider that is actively pursuing AI-enabled business models and had more than fifty AI agents in operation in early 2026, to illustrate how companies build capabilities for allocating decision rights between humans and AI at scale.<sup>3</sup>

## NOT ALL AI DECISIONS ARE EQUAL

A key to effective AI adoption is recognizing that not all decisions are alike. Some are predictable and low risk, while others are ambiguous and consequential, requiring different approaches to automation and oversight. The same AI capability can be safe in one context and risky in another.

The AI Decision Matrix addresses this challenge using two dimensions: ambiguity and risk (see the figure).

1 Agentic AI systems can plan, act, and invoke tools to execute multistep tasks with limited human involvement, enabling new business models.

2 In 2025–2026, we conducted 30 interviews with 27 executives at nine global companies as part of MIT CISR research projects on business processes, AI-enabled business models, and decision rights in the agentic AI enterprise. Companies represented telecommunications, financial services, manufacturing, and law. Interviewees included CIOs, chief AI and data officers, chief governance officers, heads of strategy and knowledge management, and other executives on these leaders' teams.

3 The One New Zealand content in this briefing is based on P. Weill, I. M. Sebastian, S. L. Woerner, and G. Benedict, "Business Models in the AI Era," MIT CISR Research Briefing, Vol. XXV, No. 10, October 2025; P. Weill and S. Collins, "Interview with Summer Collins, Chief AI and Data Director of One NZ," February 25, 2026, MIT Center for Information Systems Research, video, 16:38; and material One NZ approved for MIT CISR's use.

- **Ambiguity** describes how clearly data determines the answer. Low-ambiguity decisions are repeatable and predictable. High-ambiguity decisions allow multiple reasonable (and sometimes unreasonable) interpretations.

- **Risk** describes the consequences of getting a decision wrong. Low-risk decisions have little impact if wrong and are reversible. High-risk decisions have significant financial, operational, or reputational impact if wrong and are often harder to reverse.

Combined, these dimensions distinguish four types of decisions—routine, consequential, exploratory, and strategic—and clarify how humans and AI should be involved.

### Routine Decisions (Low Ambiguity, Low Risk)

Routine decisions are well defined and have limited consequences if wrong, making them suitable to automate. An example of a routine decision at One NZ is creating audience segments from the customer base for marketing campaigns. A customer segmentation agent translates marketer requests into SQL queries, reducing audience creation time by sixty percent. AI can handle most of the work because the decision logic and data requirements are clear and the risk is low.

### Consequential Decisions (Low Ambiguity, High Risk)

Consequential decisions are well-defined, but errors are costly. An example at One NZ is determining whether an AI agent can act on behalf of customers in customer service, such as upgrading plans or opening support tickets. AI agents resolve more than 60 percent of One NZ's in-app questions. The company determined that clear guardrails enable automation in these cases, while relying on a human to respond to every question would slow service without necessarily improving outcomes.

### Exploratory Decisions (High Ambiguity, Low Risk)

Exploratory decisions involve interpretation, creativity, or uncertainty, but consequences of errors are limited. An example at One NZ is using AI to generate marketing campaign ideas

and content. Marketing teams evaluate which outputs fit the brand, customers, and campaign goals. AI expands options and accelerates experimentation, but humans can revise mistakes before they become consequential.

### Strategic Decisions (High Ambiguity, High Risk)

Strategic decisions involve both uncertainty and significant consequences. An example at One NZ is deciding how to balance demand and capacity in the network. A power supervisor AI agent analyzes network data and recommends how to manage power usage, but the decision is human-led because it involves tradeoffs among service reliability, customer experience, cost, and infrastructure priorities. One NZ's Chief AI and Data Officer Summer Collins said, "We use AI to, in essence, stand up a team of what would have probably been thirty people to help us make better decisions." Here, AI supports leadership judgment by integrating data and surfacing options.

## DECISION-MAKING INVOLVES FRAMING, ACTING, AND LEARNING

A decision is not a single moment of choice. It is a set of activities through which a company interprets a situation, takes action, and learns from outcomes. We break a decision into three components—frame, act, and learn—that illustrate different ways AI and humans participate in that decision:<sup>4</sup>

- **Frame:** define the problem, assumptions, stakeholders, constraints, and success criteria
- **Act:** gather information, generate and evaluate options, recommend or authorize action, and execute approved steps
- **Learn:** monitor and interpret outcomes, ensure accountability for outcomes, and update the decision system over time

Together, these components clarify how the company defines problems, executes decisions, remains accountable for results, and adapts over time.

## DECISION RIGHTS SHIFT WITH RISK AND AMBIGUITY

As decisions vary in ambiguity and risk, companies need to adjust how they allocate activities between humans and AI to frame and act on decisions and learn from their outcomes, while maintaining clear human accountability. Each type of decision requires a different approach:

**Routine decisions: Automate.** Companies can codify framing in advance for these decisions. They automate acting and learning with a human process owner monitoring patterns. Many companies keep humans closely involved in decisions while building confidence with AI agents.

For example, marketing teams at One NZ review AI-generated audiences before launching campaigns. Over time, as teams gain confidence in AI performance, they may give it more autonomy.

**Consequential decisions: Manage.** Companies must design these decisions for reliability. Framing establishes risk tolerance, tradeoffs, and guardrails. Humans monitor acting and manage exceptions and issues that arise. The company assigns learning to a human risk owner with clear authority to improve the AI and make other modifications when needed.

At One NZ, leaders established the company's risk tolerance. One non-negotiable constraint was that they could never get a price wrong for a customer. Based on this constraint, One NZ determines which decisions to automate in customer interactions and which require a human in the loop. The company introduces agents gradually, starting with small cohorts to validate outcomes before scaling.

**Exploratory decisions: Enable.** These decisions involve experimenting with AI and require flexibility. Framing evolves through interaction between humans and AI. Humans monitor acting. Learning remains with human individuals or teams using AI. Leadership must ensure governance keeps pace as experiments transition into higher-risk applications.

One NZ marketing teams began using agents for content creation tasks, framing campaign goals with and without the help of the AI agents, delegating content creation to them, and learning from customer responses. The teams are now developing agent-led end-to-end autonomous campaigns with adaptation based on customer behavior. This shift from exploratory decisions, which are low risk, to strategic decisions, which are high risk, raises new questions about governance and accountability.

**Strategic decisions: Lead.** Strategic decisions require strong human oversight. Human executive sponsors own framing and learning, and AI supports acting.

To optimize the network and improve resilience in the event of power outages, One NZ deployed fifteen task-based agents to analyze network data. Orchestration agents manage the task-based agents, integrating insights and supporting human decision-making with increasingly better data and analysis. The company has introduced progressively more closed-loop automation across the network agents.

<sup>4</sup> These components draw on multiple academic research streams on individual judgment, decision processes, and organizational design, with each stream informing different aspects of how participants in decisions frame, act on, and learn from them.

## ENTERPRISE CAPABILITIES SUPPORT DECISION RIGHTS

The AI Decision Matrix clarifies how decision rights should shift with risk and ambiguity. One NZ has made organizational design choices to ensure it accommodates these shifts consistently. Business units own the framing of each AI use case, anchoring agents in business need, value, and domain expertise. The company’s center of excellence for AI and data owns AI and data foundations, architecture, and delivery. A unified data platform provides access to structured and unstructured data, while a horizontal orchestration layer coordinates across domain specific agentic platforms. One NZ’s Responsible AI Policy,<sup>5</sup> training,<sup>6</sup> guardrails embedded in AI solution lifecycle gates through a standardized “Trust Trail,”<sup>7</sup> and iterative rollout of AI agents to small cohorts support aligned decision-making.

One NZ views AI agents as colleagues requiring ongoing management. As Collins noted, “AI agents often start as very enthusiastic interns” that need coaching, guardrails, and the

right data. No AI agent is deployed without a named business owner accountable for its outcomes and learning. Owners monitor performance, refine data, test accuracy, and improve agents over time. For a financial reconciliation agent—a Consequential case in the matrix—the company’s chief financial officer was its accountable business owner, while a trusted, AI-savvy finance leader embedded in the AI and data team developed the agent and handled day-to-day stewardship with solution and data engineers. This reduced coordination overhead and accelerated progress.

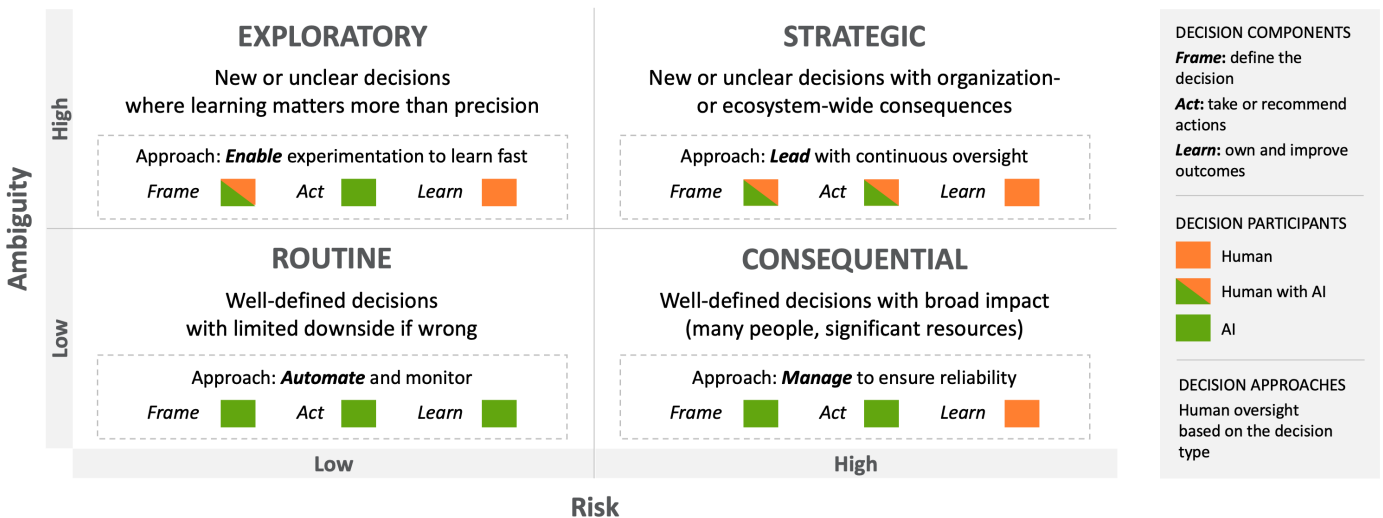
## MANAGE AI AS A PORTFOLIO OF BUSINESS DECISIONS

As the company embeds AI in core processes, leaders need to look beyond the AI use case to the business decision it affects. The AI Decision Matrix helps them assess each use case by asking how ambiguous the decision is and what could happen if the decision is wrong. It also outlines how humans and AI participate in framing the issue, taking action, and learning from outcomes.

We predict the companies that get this right will move faster, reduce risk, and build trust in how they use AI. Mastering routine and consequential decisions can deliver near-term gains in efficiency and scale. Designing for exploratory and strategic decisions can unlock longer-term value through innovation and better judgment. The important shift is from asking where AI can be used to managing the business decisions it shapes.

5 One New Zealand, AI Trust Report: New Zealanders’ Attitudes Towards AI in 2025, October 20, 2025, <https://content.one.nz/6f/46/388ce4174a308072737a742bf4b4/ai-trust-report-2025-one-nz.pdf>.  
 6 “Responsible AI at One NZ,” Media, One NZ, October 16, 2025, <https://media.one.nz/katikarai>.  
 7 One NZ’s Trust Trail is One NZ’s framework for navigating AI risk, ensuring compliance, and embedding responsible AI at scale.

Figure: The AI Decision Matrix



Source: 29 interviews with 27 executives at nine global organizations in 2025–2026 as part of MIT CISR research projects on business processes, AI-enabled business models, and decision rights in the agentic AI enterprise. Companies represented telecommunications, financial services, manufacturing, and law. Interviewees included CIOs, chief AI and data officers, chief governance officers, heads of strategy and knowledge management, and other executives on these leaders’ teams.

## MIT CENTER FOR INFORMATION SYSTEMS RESEARCH (CISR)

MIT CISR helps executives meet the challenge of leading increasingly digital and data-driven organizations. We provide insights on how organizations effectively realize value from approaches such as digital business transformation, data monetization, business ecosystems, and the digital workplace. Founded in 1974 and grounded in MIT's tradition of combining academic knowledge and practical purpose, we work directly with digital leaders, executives, and boards to develop our insights. Our research is funded by member organizations that support our work and participate in our consortium.

### CISR RESEARCH PATRONS

AlixPartners  
Avanade  
Cognizant  
Collibra  
IFS  
PwC

### CISR SPONSORS

ABN Group (Australia)  
Alcon Vision  
ANZ Banking Group (Australia)  
AustralianSuper  
Banco Bradesco S.A. (Brazil)  
Barclays (UK)  
BNP Paribas (France)  
Bupa (Australia)  
CalSTRS  
Caterpillar, Inc.  
Cemex (Mexico)  
Cencora  
CIBC (Canada)  
Commonwealth Superannuation Corp. (Australia)  
Cuscal Limited (Australia)  
Dawn Foods  
DBS Bank Ltd. (Singapore)  
Doosan Corporation (Korea)  
Ericsson (Sweden)  
Fidelity Investments

Fomento Economico Mexicano, S.A.B., de C.V.  
Genentech  
HCF (Australia)  
Hunter Water (Australia)  
International Motors  
JERA Co., Inc.  
JPMorgan Chase  
Kaiser Permanente  
Keurig Dr Pepper  
Mallesons (Australia)  
Mater Private Hospital (Ireland)  
National Australia Bank, Ltd.  
Nomura Holdings, Inc. (Japan)  
Nomura Research Institute, Ltd. Systems Consulting Division (Japan)  
Novo Nordisk A/S (Denmark)  
OCP Group  
Pentagon Federal Credit Union  
Principal Life Insurance Company

Ralliant  
Reserve Bank of Australia  
RTX  
Saint-Gobain  
Scentre Group Limited (Australia)  
Schneider Electric Industries SAS (France)  
TabCorp Holdings (Australia)  
Telstra Limited (Australia)  
Terumo Corporation (Japan)  
Uniting (Australia)  
Vanguard  
WestRock Company  
Xenco Medical  
Zoetis Services LLC

MIT CISR is funded by our members, and we gratefully acknowledge their financial support and their many contributions to our work.

Membership and benefits: [c isr.mit.edu/content/member-benefits](https://c isr.mit.edu/content/member-benefits)

MIT CISR research publications: [c isr.mit.edu/research-library](https://c isr.mit.edu/research-library)



**MIT Sloan School of Management**  
Center for Information Systems Research

One Main Street, E90-9th Floor  
Cambridge, MA 02142

t 617-253-2348 | e [c isr@mit.edu](mailto:c isr@mit.edu)

[c isr.mit.edu](https://c isr.mit.edu) | 

**Team** | Isobela Byerly-Chapman, Margherita Di Pinto, Christine G. Foglia *Associate Director*, Dorothea Gray-Papastathis, Cheryl A. Miller, Ina M. Sebastian, Alan Thorogood, Nick van der Meulen, Peter Weill *Chairman*, Barbara H. Wixom, Stephanie L. Woerner *Director*